DESCRIPTION

WIRELESS NETWORK SECURITY

The present invention relates to wireless communication networks and to network security.

Wireless communication networks have increased in popularity in recent years. Mobile phones, pagers, personal digital assistants, mobile computers and other communication devices all take advantage of wireless communication technologies. The use of wireless communication networks has become commonplace in business and industrial environments and increasingly these technologies are becoming more popular in a domestic setting.

The ambient environment, for example the home of the future, may be based on ad-hoc wireless networking technologies, such as ZigBee and Bluetooth. Electronic devices in the ambient environment will form part of this network. Examples could include, but are not limited to, household alarms, thermostats for central heating / air conditioning, light switches and smoke detectors. An example scenario could be one where a home owner requests to find his car keys and the keys send a reply back to the owner.

The nature of wireless communication is such that broadcasts within the domestic network need not be restricted to the vicinity of the home. In all likelihood, domestic networks will extend beyond the confines of the living environment and may overlap with a neighbouring network. This makes the domestic network vulnerable to tampering and security breaches by a third party. The third party may be an innocent user of an adjacent network or a more determined trespasser.

Security in ad-hoc domestic wireless networks is an issue of wireless communication which has not been addressed as thoroughly as other aspects of the wireless technologies. In the domestic setting, it will be necessary to avoid, or reduce, interference from neighbouring networks

or passers-by to counter information being gleaned from the network and to counter other security breaches.

Unlike conventional wired networks, where a determined trespasser must gain physical access to the wired link or exploit security weaknesses in firewalls and routers, attacks in an ad-hoc wireless network may arise at any point within the network or at any location capable of receiving the wireless broadcasts. Ad-hoc wireless networks do not have a clear line of defence and measures must be taken to counter or immunise against either innocent or malicious security breaches.

WO 02/078210 describes a wireless communication system in which parts or all of a message are transmitted over different communication paths defined by repeaters in the network, thereby preventing obstacles that inhibit wireless communication on one path from blocking an entire message. WO '210 specifically addresses the problem of fading where messages transmitted between two transceivers in a network are lost or garbled, and does not discuss security against intentional tampering or interference with messages.

The expression 'tampering' is intended to indicate all forms of intentional or malicious interference with a signal, including interference with the data content of a signal. Tampering may include effecting an attempted denial-of-service whereby information is prevented from arriving at the intended destination, e.g. by jamming, and access attempts whereby a determined trespasser intentionally attempts to gain unauthorised access to the network.

An ad-hoc wireless network is one which accepts (and rejects or sometimes drops) member devices automatically, i.e. with a minimum of fuss and without need for human intervention. Such a network may only have one device until a similarly-enabled device approaches it or is placed in its vicinity. A master-slave configuration is generally preferred in such networks. Both ZigBee and Bluetooth use such a configuration. However, peer-to-peer configurations are also used. The present invention is applicable to these and other wireless network topologies.

In the future it is envisaged that such an environment would have tens of such devices. At any one time, therefore, these devices would be starting, participating in and ending conversations with other devices. If the two conversing devices are not in each other's immediate vicinity (i.e. they are out of range of each other's wireless transmissions) then these conversations take place via intermediaries – routers or repeaters – which themselves are capable of conversing with other devices. There is a clear need for security measures in the management of these networks.

An object of the present invention is to provide improved security in wireless networks.

According to one aspect, the present invention provides a method of communication over a wireless communications network, the network comprising at least first and second transceivers, linked by wireless communication paths, each path including at least one repeater disposed within the network for the propagation of messages, the method comprising the steps of:

transmitting a plurality of signals that make up a message, through the network to the second transceiver; and

determining from received signals, whether one or more of the signals has undergone tampering.

According to another aspect, the present invention provides a receiver for receiving messages over a wireless communications network and for detecting tampering of the signals in the wireless network, comprising:

means for receiving a plurality of signals that make up a message, from the network; and

means for determining from the received signals whether one or more of the signals has undergone tampering.

According to another aspect, the present invention provides a method for detecting the presence of an unauthorised device attempting to connect to a network, comprising the steps of:

transmitting a first message, onto a network, which first message includes spurious data which purports to be data that enables or maintains connection of a device to the network; and

detecting subsequent use of that data to identify an unauthorised attempt to connect to the network.

According to another aspect, the present invention provides a device for use on a network, the device comprising:

a transmitter for transmitting a first message onto the network, which first message includes spurious data which purports to be data that enables or maintains connection of a device to the network; and

detection means for detecting subsequent use of that data, by another device, to identify an unauthorised attempt to connect to the network.

Embodiments of the present invention will now be described by way of example and with reference to the accompanying drawing in which:

Figure 1 shows a schematic representation of an ad-hoc wireless communication network according to a preferred embodiment of the present invention; and

Figure 2 shows a schematic representation of an ad-hoc wireless communication network according to a second embodiment of the present invention.

With reference to figure 1 there is shown an ad-hoc wireless communication network 10 according to an embodiment of the invention. The network 10 is preferably located within a domestic setting, although the network 10 may also be located in other environments.

The network 10 includes at least two transceivers 1, 2 which may be any suitable short-range transceiving devices capable of operating in the 2.4 GHz (ZigBee) frequency band, although other frequency bands may be used. The transceiving devices are preferably compliant with the ZigBee standards-based wireless technology, but may also be compliant with Bluetooth, 802.11 and other wireless standards.

The transceivers 1, 2 are linked by at least two wireless communication paths 4, 4'. The communication paths 4, 4' operate in the frequency bands of the transceivers 1, 2. The communication paths 4, 4' are each defined by at least one, but preferably more, different repeaters $3_1...3_n$, $3_1'...3_n'$ located within the network. The repeaters may be any suitable short-range transceiving devices capable of operating at the frequency of the transceivers 1,2 and are compliant with the preferred wireless technology standard of the transceivers 1, 2. It will be recognised that the number of repeaters $3_1...3_n$, $3_1'...3_n'$ required in the network 10 will depend on the dimensions of the domestic environment and the number of desired different communication paths 4, 4'. Therefore, the number of communication paths 4, 4' is interrelated with the number of repeaters $3_1...3_n$, $3_1'...3_n'$.

In preferred arrangements, two transceivers 1, 2 are linked by two communication paths 4, 4', each path defined by at least one different repeater $3_1...3_n$, $3_1'...3_n'$ located within the network 10. One of the transceivers 1 transmits a plurality of signals making up a message, through the network 10 to the other transceiver 2. The other transceiver 2 determines, from the received signals, whether one or more of the signals has undergone tampering and thus whether the network is maintaining its integrity (ie. it has not been breached). Each path 4, 4' may have the same number of repeaters $3_1...3_n$, $3_1'...3_n'$ or a different number of repeaters.

In one arrangement, one transceiver 1 divides the signals of a message across the two communication paths 4, 4' and transmits a portion of the signals on one path 4 and the remainder of the signals on the other

path 4'. The signals of the message may be divided into equal portions or unequal portions. The divided signals may be consecutive or randomly divided so as to be non-consecutive. The divided signals sent on the same path may be separated in time.

A significant benefit of this arrangement is that no one communication path carries all the signals of the message. Thus eavesdropping on a particular path does not provide all the information in the message and jamming prevents only a portion of the message from being received.

In another arrangement, one transceiver 1 sends signals of the message on one communication path 4 and sends the same signals on the other communication path 4'. The signals may be sent on both paths 4, 4' at the same time, or may be sent at different times, either overlapping or separated by a predetermined interval of time.

A significant benefit of this arrangement is that jamming of a particular communication path 4 does not prevent the signals of the message from being received on the other communication path 4' by the other transceiver 2.

With reference to figure 2, there is shown an alternative network configuration comprising a plurality of nodes N distributed around an ambient network 10, comprising two subnets 11, 12 each comprising a collection of nodes N, eg node 3_1 in subnet 11 and node 3_2 in subnet 12. A controller node 13 may also be provided, the function of which will be described in greater detail hereinafter. An attacker node 14 attempting to obtain illegal access to the network 10 is also illustrated. A first transceiver node 1 (acting as a message originator) in subnet 12 may communicate with a second transceiver node 2 (acting as a message recipient) over the network either directly or using any convenient set of intermediate nodes N acting as repeaters, eg. node 3_1 and 3_2.

It will be understood that any of the nodes N could act as a message originator or a message recipient or as repeater to other nodes' messages.

It will be understood that the existence of a large number of nodes in the network that can act as repeaters for two communicating transceivers results in a plurality of possible pathways between the transceivers, many of which pathways are completely independent and some of which share at least some common links between adjacent nodes.

An aspect of the invention is that tampering with transmitted signals can be detected by the receiving device. A number of ways are provided for achieving this.

Tampering may be evidenced by delay, disruption or termination in transmitted or retransmitted signals over multiple paths, or by the existence of corruption of data within a message, or the use of illegal data within a message. This illegal data could be spurious data previously deliberately propagated onto the network as will be described hereinafter.

Delays caused by tampering with signals may be detected by having prior knowledge of expected delays in normal communication paths. For example, the presence of a repeater $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ introduces a known delay into the signal propagation time. An increasing number of repeaters in a path 4, 4' will increase the signal propagation time by an amount related to the number of repeaters. The amount of time contributed by the repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ will correspond to an expected delay in the signal at the receiving transceiver 2. Knowledge of this delay may allow the receiving transceiver 2 to determine if signal tampering occurs, since a delay greater than the expected delay may be indicative of a security breach.

Deliberate delays may be introduced between signals transmitted on different paths and these delays may also be taken account of in detecting tampering.

In a preferred arrangement, the transceivers 1, 2 in the network 10 each include a signal processing module (not shown) to compare the signals received on the two communication paths 4, 4'. It will be understood however, that an attack amounting to a denial-of-service on a

particular path 4 will either result in no signals being received on that path 4 or else signals will cease at the moment jamming commences.

In one arrangement, the signal processing module will notify the receiving transceiver 2 of signal tampering if the transceiver 2 receives signals on one communication path 4 only. This condition applies to both arrangements whereby either signals of the message are divided and sent across two communication paths 4, 4' or the same signals are sent across two communications paths. A failure to receive any signals at all on one communication path 4, 4' would indicate a denial-of-service attack on that path.

In another arrangement, the signal processing module will notify the receiving transceiver 2 of signal tampering if the transceiver 2 receives signals on one communication path 4 that do not match the signals received on the other communication path 4', where the transceiver 2 was expecting to receive the same signals on both communication paths.

In another arrangement, the signal processing module will notify the receiving transceiver 2 of signal tampering if the transceiver receives signals on both communication paths 4, 4' that indicate at least one signal is missing. This condition applies to signals of a message that were divided and sent across two communication paths 4, 4'. The signal processing module is preferably adapted to combine portions of signals received on one communication path 4 with portions of signals received on the other communication path 4'. If the combined portions do not make up a whole transmitted message, an attack on the signals is indicated. Preventing one or more signals from being received at the receiving transceiver 2 amounts to a denial-of-service attack.

In the preferred arrangement, any notification by the signal processing module to the receiving transceiver 2 to the effect that signal tampering has occurred to one or more signals of the message will preferably lead the transceiver 2 to assert a security breach to the network 10. This may take the form of an alarm signal (not shown) that is transmitted across all communication paths 4, 4' to each transceiver 1,2 of

the network 10. The assertion of a security breach may inhibit further network communication, either temporarily, or indefinitely, by shutting down the communication path 4, 4' under attack (ie. partly inhibiting the network) or else shutting down all communication paths 4, 4' (ie. fully inhibiting the network).

The shutdown may be accomplished by one or more transceivers 1, 2 or repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ refusing to transmit further signals or else halting signals currently being transmitted. The shutdown may be for a prescribed time interval or until the network 10 is manually reset by an authenticated user. The shut down may comprise only isolation of a node from which there is evidence of tampering.

In another arrangement, the transceiving devices in the network 10 may each invoke a lock-down condition in response to a security breach assertion to the network 10. The lock-down condition may cause the transceiving devices to refuse any network access requests either temporarily or indefinitely, even if the access request is known to be authentic. The network 10 may only be reset thereafter using a reset authentication code.

In another arrangement, the one or more repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ defining a communication path 4, 4' that is under attack may, in response to the security breach assertion, dynamically re-route the communication path so as to physically avoid that part of the path known to be the region of the signal tampering. The repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ may re-route the transmitted signals along the next available communication path 4, 4' in a manner so as not to interfere with the transmitted signals of that path. If this is not possible, the repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ may select the next available communication path 4, 4' known to be free of transmitted signals. The re-routed communication may be re-routed to a known secure path.

In such an arrangement, the network 10 is not shut down and the receiving transceiver 2 monitors the other available communication paths 4, 4' for any re-routed signals, allowing additional time for their propagation across the re-routed path.

Preferably, a transceiver 1, 2 asserting, or receiving, an alarm signal may sound either an audible or visual alarm, or a combination of both. Of course, it is to be understood that any suitable form of alarm could be used to indicate a security breach. Alternatively, or in addition, perceived security breaches may detailed in an event log which could be used to trace the activities of an attacker.

In a preferred arrangement, at least one of the communication paths 4, 4' is configured to be deliberately insecure and vulnerable to security breaches. This may be accomplished by deliberately positioning one or more repeaters $3_1'...3_n'$, which define the communication path 4', in locations which are known to be close to the extremities of the wireless network environment. In the domestic setting, examples would be, although not limited to, an attached garden, out-house or garage. The signals transmitted on the insecure communication path 4' would be particularly vulnerable to eavesdropping and denial-of-service attacks.

The signal processing module may be adapted to particularly scrutinise the signals received on the insecure communication path 4' for evidence of tampering, or else to rapidly identify jamming of the signals on the path 4'.

In another arrangement, at least some of the signals of the message could be adapted to contain deliberately spurious or bogus information. The spurious information could be, for example, fake network authentication details, false user ids or fake personal identification numbers (PINs). The spurious information could be transmitted across the network 10 as part of the signals of the message. A significant benefit of this arrangement is that eavesdropping the network would only in part obtain authentic information. Any spurious information obtained from a spurious signal would be worthless to the determined trespasser. However, the use of spurious signals is advantageous to the security of the network, since any attempt to access the network 10, or influence transceiving devices in the network 10, using the spurious information

would be readily identified by the receiving transceiver 1,2 concerned as being fake and arising from earlier signal tampering.

It is also possible, in another arrangement, that the spurious signals could be sent separately to the signals of the message.

In another arrangement, the spurious signals could be sent only across the insecure communication path 4'. The combination of the spurious information and vulnerability of the path 4' to attack would allow innocent passers-by or determined trespassers to readily obtain information from the network 10, information which is known to the transceiving devices of the network 10 to be fake.

In other arrangements, the spurious signals may be originated by the first transceiver 1, and transmitted through the network 10, to the second transceiver 2 across some or all of the communication paths 4, 4'. Alternatively, the spurious signals may be originated by the second transceiver 2, and transmitted through the network 10, to the first transceiver 1 across some or all of the communication paths 4, 4'. Alternatively, the spurious signals may be originated by a transceiver and directed back to itself, by propagation over the network.

The spurious signals may also be originated by the one or more repeaters $3_1 \ldots 3_n$, $3_1' \ldots 3_n'$ and transmitted over some or all of the communication paths 4, 4'.

The spurious signals could be sent periodically or at random times. The spurious signals could be transmitted over a pathway that has previously been identified as having been subject to a security breach, ie. tampering has been detected. In this event, the spurious signals could be manifested as an apparent (but fake) continuation of a message that has actually been re-routed due to the alarm condition indicated by the detected tampering. In this way, an attacker node may not be made aware that its existence has been detected until after human intervention can be made.

The spurious signals could comprise bogus authentication handshakes between selected devices, in particular those located in relatively insecure parts of the network.

In order to detect tampering by the unauthorised use of spurious signals, it may be necessary for at least some authorised devices in the network 10 to be able to distinguish spurious signals (eg. bogus information) that have deliberately been used by authorised devices from spurious signals that originate from an unauthorised device or attacker attempting to gain access to the network.

One way of achieving this is to maintain Look-Up Tables (LUTs) (not shown) in the devices on the network containing spurious signal entries and their conditions for use. The tables and entries may be common throughout the network 10 or specific to certain devices. A particular transmitting device 1, 2 may select an entry at random, or in a predetermined order, from the LUT for sending in accordance with a predetermined condition of use. On receiving a signal, a receiving device, eg. transceiver 2 compares the signal to the entries in the LUT stored on the transceiver 2, checks the condition of use and thereby determines whether the spurious signal indicates tampering. The conditions for use may include timing constraints (eg. when the spurious information may be passed), association with other data, including source and destination identities, or any other condition of use that can be assessed to verify whether the spurious signal originated from an authorised device or an unauthorised device.

Rather than, or as well as, using a LUT, each transceiving device of the network 10 may generate spurious signals using a mathematical algorithm common throughout the network 10.

It will be understood that where spurious signals are propagated by a single transceiver over the network and back to itself, only that transceiver need be aware of the status of such spurious signals. Still further, for the purposes of propagation of spurious information to potential attacker devices 14, a network node need not even transmit the spurious

information to another authentic node in the network. Rather, all that is required is that the network node transmits spurious information which purports to be data that enables or maintains connection of a device to the network. Then, if that data is used by another device in an attempt to connect to the network, tampering will be detected by the network node that originated the spurious data.

In another preferred arrangement, the repeaters $3_1...3_n$, $3_1'...3_n'$ of the network 10 are preferably polled by the transceiver 1 prior to transmitting the signals of the message. By polling the repeaters in the network 10, the transceiver 1 is able to determine the availability of the repeaters. The available repeaters indicate which communication paths 4, 4' are available for the transmission of signals. The repeaters $3_1...3_n$, $3_1'...3_n'$ are preferably polled by sending an enquiry message from the transceiver 1.

Preferably, the transceiver 1 selects at least two communication paths 4, 4' based on the responses of the polled repeaters $3_1...3_n$, $3_1'...3_n'$.

In a preferred arrangement, the transceiver 1 will select repeaters $3_1'...3_n'$ known to define an insecure communication path 4', in preference to other available communication paths 4, on which to send spurious signals.

In another arrangement, the repeaters $3_1...3_n$, $3_1'...3_n'$ are adapted to inject spurious signals into the signals of the message. The repeaters may use either a LUT or an algorithm, or combination of both, in order to prepare spurious signals for sending through the network, in accordance with previous arrangements. The repeaters may send the spurious signals periodically or randomly, without interfering with the signals of the message.

With further reference to figure 2, another arrangement makes use of a central controller node 13 to manage the network security. In either master-slave or peer-to-peer configurations a controller node may be provided which is connected to a device capable of larger data storage

and processing capabilities than other nodes. This controller node could be responsible for the network security measures.

There may be several controllers in one ambient environment organised in a hierarchical manner, such that there is an overall master controller. In a peer-to-peer network configuration, the controller may be designated as a node that is queried for authentication or other security related details.

An ad-hoc network can consist of just one device. The network starts growing when another enabled device is in the vicinity and joins the network. Preferably, the initial device should be the controller node. Security breaches can begin when a second legitimate device which is in the vicinity, attempts to join the ad-hoc network. An eavesdropping device could capture the details exchanged between controller and new device. These details may be used later by an attacking device in order to access the network.

One of the proactive security measures proposed is that even at its smallest (when it consists of just the one device) the ambient network can generate spurious information. An example would be a totally fabricated conversation between two devices (the controller pretends to be both message originator node 1 and message recipient node 2). An eavesdropper picking up these transmissions and attempting at a later time to use the information gleaned in order to access the network would then be identified. In this arrangement, only the controller 13 originates spurious data (and controls initiation of conversation), therefore the controller easily recognises an attacker 14 attempting to use this spurious data. In another arrangement two different controllers could masquerade as nodes 1 and 2.

This attack pattern (and proactive security measures) can be scaled-up no matter what the size of the ambient network.

Although the described embodiments are ideal for improving security in ad-hoc home wireless communication networks, it will be

recognised that the principle can be extended to other types of wireless communication network e.g. non-domestic wireless networks.

Other embodiments are intentionally within the scope of the appended claims.